arXiv:1608.03608v1 [cs.SE] 11 Aug 2016

Aristotle vs. Ringelmann: A response to Scholtes et al. on Superlinear Production in Open Source Software

Thomas Maillart*

UC Berkeley, Berkeley, USA

Didier Sornette†

ETH Zurich, Zurich, Switzerland

(Dated: August 15, 2016)

On December 19, 2015, Scholtes et al.'s work was published online in the Journal of Empirical Software Engineering [1], in which they challenged the exciting findings that we (with another co-author) presented in 2014, showing that open source software production exhibits superlinear productive bursts [2]. We presented our findings as the first quantification of Aristotle's famous adage: *The whole is more than the sum of its parts*. In contrast, Scholtes et al. referred to Maximilien Ringelmann, a French agricultural engineer (1861-1931) who discovered the tendency for individual members of a group to become increasingly less productive as the size of their group increases [3]. Since Ringelmann, the topic of collective intelligence has interested numbers of researchers in social sciences, as well as practitioners in management aiming at improving the performance of their team. And indeed, in most research and practice case studies, the Ringelmann effect has been found to hold. Our results challenge common wisdom, precisely because the Ringelmann effect has been replicated in countless experiments. We are glad that the study by Scholtes et al. allows us to explain further our findings by highlighting the methodological aspects that condition our diverging results. In particular, the definitions of group sizes and measurement periods are essential to explain the different conclusions. This healthy debate enriches the understanding on group productivity, in particular in the case of open source software production.

*Electronic address: maillart@berkeley.edu

†Electronic address: dsornette@ethz.ch

I. PRODUCTIVITY & TEAM SIZE

We start by a fundamental conceptual remark that illuminates one key difference between the approach of Scholtes et al. [1] and our's [2]. Scholtes et al. consider production in the mean, using as metric the *"average output per team member"* (Introduction, 2nd paragraph, line 3), and argue that it increases when synergy effects are present and decreases due to *"communication and coordination overhead"* (which surges with larger teams). In contrast, we argued and demonstrated [2] that using an average output is misleading if not plain wrong in the presence of a highly bursty dynamics characterised by power law tail distributions with small tail exponents. Indeed, we reported that the distribution of contributions per developer is heavy-tailed. This empirical fact is also cited by Scholtes et al. and well-documented in the open source software production literature and for other open collaboration projects, such as wiki's and Wikipedia [4–7]. In open collaboration, a few contributors account for a majority of the performed work, whether it is counted in lines of code, commits, files modified, and so on. This is one of the features associated with the fact that the distribution of contributions, counted in commits or in lines of codes, possesses a power law tail of the form $P(X > x) \sim 1/x^{\mu}$ with $\mu < 1$ [2]. Such distributions are *wild* [8] in the sense that their two first statistical moments (mean and variance) are undefined and diverge as the sample grows. For such heavy-tail distributions, reasoning in mean is fundamentally erroneous. For a finite number n of developers in the project, it is easy to show that the average production scales as $\sim n^{1/\mu}$ for $\mu < 1$ and $\sim n$ for $\mu \geq 1$. Defining productivity as the ratio of the total production by the number n of team members, this shows that productivity scales as $\sim \frac{n^{1/\mu}}{n} = n^{\frac{1}{\mu}-1}$ for $\mu < 1$ and is constant for $\mu \geq 1$. This latter case is the null hypothesis of an approximate constant output per team member. Superlinear production is quantified by $\mu < 1$, leading to a growing productivity per team member, the larger the team. Searching for a superlinear productivity is different from seeking a superlinear production, the former requiring $\frac{1}{\mu} - 1 > 1$, i.e., $\mu < 1/2$, while the later just needs $\mu < 1$. In our dataset of 168 projects, we found that only four projects are characterised by $\mu < 0.5$ (while most obey $\mu < 1$). Scholtes et al. have used a dataset of 58 projects, which implies a $\approx 0.8\%$ chance of finding one project where the average productivity scales superlinearly with the team size. But, as mentioned, the novel non-trivial discovery that we previously reported [2] deals with the superlinear production, and not superlinear productivity.

More generally, the definition of productivity needs to be carefully addressed. Indeed, an open source software community does not come into being fully grown. It starts rather small and then grows progressively with the project. When growing, the community bears increasing communication and coordination costs as pointed out by Scholtes et al. While recognizing the importance of different team sizes, Scholtes et al., decided to pick projects meeting the following criteria: (i) at least one year of activity, (ii) 50 different active developers, and (iii) being among the 100 most popular projects, as measured by the number of forks on GitHub, a leading online service for open source software production. In contrast, we have carefully chosen a representative sample of the open source ecosystem with 138 projects with less than 50 developers and 30 projects with more than 50 developers. The representative sampling of projects (see Figure 1 in [2]) showed that the superlinear production is valid only for projects of sizes no more than 30 to 50 members who are active at a given time. We found statistically significant evidence that the superlinear production tends to fade away to just linear production (i.e., constant productivity per developer) for projects with more than 50 developers (see Figure 8 in [2]). In other words, the sample selection made by Scholtes et al. is heavily biased towards large projects, which are indeed exposed to more communication and coordination costs, and also exhibit less synergy effects. Our results may thus be less opposed as claimed by Scholtes et al. as they refer to different team sizes as well as different metrics.

More specifically, Scholtes et al. defined a team as the set of developers who are active at least once within a time window of 295 days, determined by the 90th quantile of the distribution of times between two consecutive commits by the same person. This definition excludes developers with a unique contribution, who nevertheless account for 40% of all contributions, as acknowledged by Scholtes et al. in Section 3.2 (end of second paragraph) of [1]. In line with our above remarks concerning the heavy-tailed distribution of contribution sizes, this definition amounts to *throw away the baby with the bath*, since it is fundamentally ill-suited to account for the fact that a few, often most senior, developers may not contribute for years in between two commits (see Figure 2 in [9]), while at the same time they may account for most of the contribution production. The definition of Scholtes et al. is thus biased with respect to the special nature of the open source software community, which is – almost by essence – different from a corporate organization, as documented in a number of management science articles (see e.g., [10] and references therein).

We nevertheless recognize that productivity may be defined in a variety of ways, each with their advantages and shortcomings. Scholtes et al. considered productivity as production per team (as an aggregate of working developers in large – 295 days – time windows), while we considered productivity as production per developer and per time unit (i.e., over a short time period of 5 days). Even though not perfect, we believe that our definition of productivity is more fine-grained than the one proposed by Scholtes et al. As such, it precisely allows capturing the subtle bursts of activity we have documented. These bursts cannot be observed by averaging developer engagement (over a team aggregate and over time).

II. COMMITS & SUPERLINEAR SCALING

Productivity is the ratio of an output (production per worker) and an input (resources allocated to the production). So far, we have mainly discussed the input, i.e., the human capital. Scholtes et al. raised concerns about the output, and claimed that the number of commits is an *erroneous* measure of production. For that, they bring forth the following argument: *"the total number of commits contributed by n developers active in a given time period cannot - by definition - be less than n, which is why the total number of commits must scale at least linearly with team size."* This apparently common-sensical claim is incorrect as we demonstrate here.

Let us consider n developers. The largest contributor makes N commits (resp. lines of code). The second one contributes $N/2^\alpha$ commits. The third one contributes $N/3^\alpha$ commits, and the n-th one contributes N/n^α commits. If $0 < \alpha < 1$ and n and N are such that $N/n^\alpha \geqslant 1$ (i.e., $n \leqslant N^{1/\alpha}$), then the total number of commits contributed by n developers is given by

$$S(n) = \frac{N}{1^\alpha} + \frac{N}{2^\alpha} + \frac{N}{3^\alpha} + \cdots + \frac{N}{j^\alpha} + \cdots + \frac{N}{n^\alpha} \sim N \cdot n^{1-\alpha}. \tag{1}$$

Thus, in this example, the total contribution of these developers grows SUB-LINEARLY as a function of the size n of the group, with exponent $1-\alpha$. Let us illustrate this demonstration by a numerical example, showing that the sublinear effect is clearly visible even for small team sizes. Let us assume that $N = 10$ and $\alpha = 1/2$. For $n = 5$ developers, the total number of commits is 32. For $n = 25$, the total contribution is equal to 86 commits, which

is 2.7 times that for the team of 5 developers (and not 5 times more). Note that for the team of 25 developers, the first contributor makes 10 commits and the last one contributes 2 commits.

We believe Scholtes et al. made a very common confusion between absolute numbers and scaling properties. More generally, in the field of fractals, this error is also often found in the literature that confuses the fact that the fractal dimension (here, the scaling exponent) tells nothing (or very little) about the density (here, the commits).

Dismissing commits as a measure of production, Scholtes et al. used the Levenshtein edit distance [11] of diffs. The Levenshtein edit distance counts the number of permutations, additions and deletions of characters necessary to match two different strings. We welcome the effort to go for more detailed measures of production, but we believe this move does not fundamentally change the picture both quantitatively and qualitatively. In our study, in addition to commits, we have further verified that superlinear scaling production holds as well for lines of code (see Figure 3 in [2]), including added and deleted lines of code taken separately (not reported). Moreover, at a qualitative level, we should stress that using the more detailed Levenshtein edit distance is not without its own problems. One may indeed argue that changing one character or a single line of code in a piece of software may be a tremendous output reflecting a major commitment in terms of human capital (think e.g., of a small edit correcting a security vulnerability) [12]. We suggest that a truly faithful measure of input would be the time effectively spent in front of a computer by a contributor in order to achieve a task for the focal open source software project. Unfortunately, this information is not nearly available to open source software researchers and, even if it would be available, one could endlessly debate on a broad (resp. narrow) definition of time, and whether the coffee break and the ping-pong sessions are actually parts of the production time: As we know, nearly all Silicon Valley software companies would include this time as truly productive time.

III. ENGAGING THE CRITICAL CASCADE MECHANISM

We appreciate that Scholtes et al. [1] engaged our work [2] and tried to replicate and falsify it. We are however disappointed that they have not addressed the very important question of the underlying critical cascade mechanism. The open source software projects

that exhibit superlinear production are also those which are propelled by self-sustained contribution cascades. In a nutshell, in these *ripe* projects, one contribution triggers on average one follow-up contribution, by the same person, or by another person in the community. One may compare this mechanism to other critical phenomena [13], to epidemic processes and to branching processes at criticality, such as for instance during a self-sustained nuclear reaction in a nuclear plant. It is truly amazing and surprising that we could find so many OSS projects that *sing and work* with such superlinear production.

Our research certainly deserves further validation and we truly appreciate a fair challenge. However, openly claiming we are wrong by completely redefining the assumptions appears unfair and misleading. More generally, we would appreciate that researchers, who try falsifying our work, stick to the obvious standards of reproducibility in science, by (i) studying in depth our claims, (ii) taking the same or similar data, and (iii) explaining how their method differ from ours by providing compelling arguments for changing working assumptions. We believe that, if Scholtes et al. would have followed these rules, they would have found results very similar to our's, but likely with some interesting differences worth exploring further.

---

[1] Scholtes, I., Mavrodiev, P. and Schweitzer, F. , From Aristotle to Ringelmann: a large-scale analysis of team productivity and coordination in Open Source Software projects, *Empir. Software Eng.*, **21**, 642-683 (2016).

[2] Sornette, D., Maillart, T. and Ghezzi, G., How Much Is the Whole Really More than the Sum of Its Parts? 1? 1= 2.5: Superlinear Productivity in Collective Group Actions. Plos one **9.8**, e103023 (2014).

[3] Ringlemann, M., Recherches sur les moteurs animés: Travail de l'homme, *Annales de l'Institut National Agronomique.* **12**1, (1913).

[4] Robles, G., Koch, S., and Gonzalez-Barahona J.M., Remote analysis and measurement of libre software systems by means of the CVSAnalY tool. *In Proceedings of the 2nd ICSE Workshop on Remote Analysis and Measurement of Software Systems (RAMSS)* 51-55 (2004).

[5] Hindle, A., German, D. M., and Holt, R., What do large commits tell us?: a taxonomical

study of large commits. *In Proceedings of the 2008 international working conference on Mining software repositories*, 99-108 (2008).

[6] Alali, A., Kagdi, H., and Maletic, J. I., What's a typical commit? A characterization of open source software repositories. *IIn: The 16th IEEE International Conference on Program Comprehension. ICPC 2008, pp 182191.* (2008).

[7] Arafat, O., and Riehle, D., The commit size distribution of open source software, In: 42nd Hawaii International Conference on System Sciences HICSS'09, 1-8, (2009).

[8] Mandelbrot, B., and Taleb, NN.. Mild vs. wild randomness: focusing on risks that matter. *Erscheint in: Diebold, F* (2007).

[9] Saichev, A., Maillart, T., and Sornette, D., Hierarchy of temporal responses of multivariate self-excited epidemic processes. *The European Physical Journal B*, **86**(4), 1-19 (2013).

[10] Von Krogh, G., Haefliger, S., Spaeth, S., and Wallin, M. W., Carrots and rainbows: Motivation and social practice in open source software development. *Mis Quarterly*, 36(2), 649-676 (2012).

[11] Levenshtein V.I., Binary codes capable of correcting deletions, insertions and reversals. In: *Soviet physics doklady*, **10**, p 707 (1966).

[12] Kuypers, M.A., Maillart, T. and Paté-Cornell, E., An Empirical Analysis of Cyber Security Incidents at a Large Organization. *Stanford Working Paper* (2016).

[13] Sornette, D., Critical Phenomena in Natural Sciences: Chaos, Fractals, Selforganization and Disorder: Concepts and Tools *Springer Series in Synergetics, Springer*, 2nd ed. (2006).
Sornette, D., Critical phenomena in natural sciences: chaos, fractals, self-organization and disorder: concepts and tools. *Springer Science & Business Media* (2006).